

A Block Circulant Embedding Method for Simulation of Stationary Gaussian Random Fields on Block-regular Grids

M. Park* and M.V. Tretyakov*

April 10, 2015

Abstract

We propose a new method for sampling from stationary Gaussian random field on a grid which is not regular but has a regular block structure which is often the case in applications. The introduced block circulant embedding method (BCEM) can outperform the classical circulant embedding method (CEM) which requires a regularization of the irregular grid before its application. Comparison of BCEM vs CEM is performed on some typical model problems.

Keywords: Stationary Gaussian random field, irregular grids, sampling techniques, circulant embedding method, symmetric block-Toeplitz matrices, block fast Fourier transform.

1 Introduction

Uncertainties are often modeled using stationary Gaussian fields [4, 10, 16, 18, 24]. Efficient generation of samples from stationary Gaussian fields is crucial for using Monte Carlo techniques, which are the backbone of uncertainty quantification simulations, in studying behavior of systems subject to uncertainties. There are a few numerical techniques for sampling Gaussian random fields on a grid. For instance, one can find a square-root of the corresponding covariance matrix using Cholesky's decomposition and then multiply the square-root by a vector of independent Gaussian random variables to simulate a sample. This is an exact method but it is rarely used in applications due to the high cost of Cholesky's decomposition in high dimensions. Another possibility is the Karhunen-Loeve expansion (see, e.g. [11, 15]), which requires knowledge of eigenvalues and eigenfunctions of the covariance operator for the Gaussian random field. In many cases of practical interest the eigenvalue problem has to be solved numerically which can be expensive, especially when eigenvalues decay slowly. Also, this method is not exact. The fast and exact method of generating large samples from stationary Gaussian fields on regular grids is the circulant embedding method (CEM) [22, 6, 23] which is widely used in various uncertainty quantification (UQ) applications such as groundwater flow simulation [11, 17], weather field forecasting [10], and liquid composite molding processes [21]. The two main drawbacks of CEM are (i) the requirement imposed on the grid to be regular while irregular grids of a block structure naturally appear in many applications (see three typical examples below) and (ii) the need to deal with non-positive definiteness of circular embedding matrices which often occur in practical applications. The remedies for the latter were considered in [19, 12, 14], here we deal with the first deficiency of CEM. To this end, we propose a new block circulant embedding method (BCEM). Let us clarify the matter using the following three examples which come from sampling a random permeability field in groundwater flow simulations.

Example 1.1. *Triangular finite element with a quadrature point located at the barycentre of the triangle.*

Consider generation of a stationary log-normal random permeability field to be used in simulations based on triangular finite elements and the Gaussian quadrature rule of degree 1 within a rectangular domain.

*School of Mathematical Sciences, University of Nottingham, Nottingham, NG7 2RD, UK (emails: min.park@nottingham.ac.uk and Michael.Tretyakov@nottingham.ac.uk).




Figure 1.1: Left: locations of nodes in 2D triangular elements. Right: regular grid nodes (gray) which contain all the original nodes (black). Note that the nodes on the right and top sides of the rectangles belong to the neighboring elements.

dimension	BCEM	CEM	CEM/BCEM
1	$(2+1)n$	$4n$	1.3
2	$(2^2 + 1)n$	4^2n	3.2
3	$(2^3 + 1)n$	4^3n	7.1

Table 1.1: Comparison of the number of nodes needed by BCEM and CEM in Example 1.2.

Assume that the rectangular domain consists of small rectangles (see Figure 1.1) and that there is no overlap of these rectangles. To perform the finite element simulation, it is sufficient to have sampled values of the permeability field at the quadrature points only (see the black circles in Fig. 1.1). The covariance matrix of the corresponding stationary Gaussian random field at all quadrature (black) points is symmetric block-Toeplitz, but the blocks themselves are not symmetric. Hence, in order for the standard CEM to be applicable, 7 extra (i.e., artificial from the point of view of sampling permeability values sufficient for the finite element simulation) points should be added to each rectangle (the gray circles in Figure 1.1 are regular grid points involving black circles). In contrast to CEM, the new method - BCEM - allows to sample values at the required points (black circles) without adding extra nodes to the grid and it does so in a very efficient way as we will see in the next sections.

Example 1.2. *Cell-centered finite volume discretization in multilevel Monte Carlo (MLMC) computation.* The multilevel Monte Carlo (MLMC) method is a Monte Carlo technique, which can give a substantial reduction of computational complexity in comparison with the standard Monte Carlo method thanks to making use of a hierarchical sampling [2, 1]. In the MLMC algorithm, when computing the difference of quantities on two consecutive grids with mesh sizes h and $2h$, the pair of fine and coarse random samples must come from the same realization of the random field. In the cell-centered finite volume discretization, which uses permeability values at the centers of cells, the locations of coarse random field do not coincide with nodes on the fine grid (see Fig. 1.2). In this case there exists a uniform grid with the mesh size $h/2$ containing both fine and coarse points, and hence it is possible to generate the required pair from the same realization by applying CEM on this finer uniform grid (grey circles in Fig. 1.2). However, this leads to an increase of both simulation time and memory requirements and, hence, to deterioration of the MLMC performance. Table 1.1 compares the number of nodes at which the random field actually needed to be sampled for MLMC (which will be the same as the number of nodes used in BCEM) against the ones on the fine, regularized grid required by CEM. The portion of unused values grows as dimension increases. The benefit of BCEM is that exploiting the block-regular structure of grids used in MLMC, it allows us to sample at the points used in finite volume simulation without need to regularize the grid by adding extra points,




Figure 1.2: Left: location of sampling points on the fine grid (black) of size h and the coarse grid (hollow) of size $2h$ using the cell-centered finite volume discretization in 2D. Right: uniform grid (gray) which contains both black and hollow points. Note that the nodes on the right and top sides of the rectangles belong to the neighboring elements.

which can result in substantial savings of both computational time and memory in comparison with applying CEM.

Example 1.3. *Conditional random field generation on block regular grids.*

The conditional random field generation based on CEM was considered in [5]. In this approach one builds a symmetric matrix of the form

$$R = \begin{bmatrix} R_{11} & R_{12} \\ R_{21} & R_{22} \end{bmatrix}, \quad (1.1)$$

where $R_{11} \in \mathbb{R}^{n_1 \times n_1}$ is a (block) circulant matrix, and $R_{22} \in \mathbb{R}^{n_2 \times n_2}$ is a covariance matrix of filed values at locations of measurements, and generate random vectors using its square root

$$R^{1/2} = \begin{bmatrix} \frac{1}{\sqrt{n_1}} F \Lambda^{1/2} & 0 \\ K & L \end{bmatrix}, \quad (1.2)$$

where $K = \sqrt{n_1} R_{21} F \Lambda^{-1/2}$ and L is a matrix such that $LL^T = R_{22} - KK^H$. Here F denotes a discrete Fourier transform matrix, and Λ is a diagonal matrix whose entries are eigenvalues of R_{11} . The computational costs of forming Λ , K^H and KK^H are $\mathcal{O}(n_1 \log n_1)$ flops, $\mathcal{O}(n_2 n_1 \log n_1)$ flops, and $\mathcal{O}(n_2^2 n_1)$ flops, respectively.

BCEM can also be used for generation of random fields conditioned on observations. As with the unconditional sampling discussed above, applications of conditional sampling often deal with grids which are not regular but have a regular block structure (see, e.g. conditional MLMC simulation in [17]). Since BCCEM requires smaller n_1 value as it does in the unconditional case, BCCEM in the conditional random field setting can outperform CEM.

BCCEM also has the remarkable feature that it is parallelizable in contrast to the standard CEM which is a serial algorithm (of course, CEM can exploit parallelism of the fast Fourier transform (FFT) but BCCEM's main ingredient is also FFT and it can benefit from FFT parallelism as well), i.e., BCCEM has a further significant advantage over CEM.

The rest of the paper is organized as follows. In Section 2 we illustrate the idea of BCCEM in the case of one-dimensional space. In Section 3 we present a multi-dimensional BCCEM. Computational complexity of BCCEM is discussed in Section 4, where some numerical experiments comparing BCCEM with the standard CEM show that already in 2D BCCEM can be three time faster in sample generation than CEM.

2 Illustration of the idea

To illustrate the idea of BCEM, we start with presenting it in the 1D case.




Figure 2.1: 1D uniform grid $\Omega_r = \{x_0, \dots, x_N\} \in \Omega = [x_0, x_N]$. Black circles represent the locations of points $s_j^{(i)} \in \Omega_s$, and the combination of black and grey circles correspond to the uniform grid $\tilde{\Omega}_s$.

Consider a uniform grid $\Omega_r = \{x_0, \dots, x_N\}$ on the interval $\Omega = [x_0, x_N]$ with a grid size $h = (x_N - x_0)/N$, and sets of points $S_i = \{s_1^{(i)}, \dots, s_\ell^{(i)}\} \subset \Omega_i = [x_i, x_{i+1}]$ with $s_j^{(i)} = x_i + \delta_j$, where $0 \leq \delta_1 < \delta_2 < \dots < \delta_\ell < h$ (see Figure 2.1). Note that δ_j are independent of the index i . The grid $\Omega_s := \bigcup_{i=0}^{N-1} S_i$ is, in general, non-uniform (it is uniform if $\ell = 1$) but it is block-uniform, i.e., the distribution of points in each sub-interval (in other words, block) Ω_i is the same.

Let $Z(x)$, $x \in \mathbb{R}$, be a stationary Gaussian random field with zero mean and covariance function $r(x)$. Our aim is to sample from $Z(x)$ on the grid Ω_s . If Ω_s is not a grid of equispaced points, then the covariance matrix of the field $Z(x)$ on Ω_s is not Toeplitz. In this case the standard CEM [22, 6, 23] cannot be applied to this covariance matrix in order to perform highly efficient computing of its square-root with subsequent generation of the required Gaussian field samples. The simplest remedy is to extend the non-uniform grid Ω_s to the uniform grid $\tilde{\Omega}_s$ by adding points (see Figure 2.1) and then apply the standard circulant embedding method, but this approach results in a substantial increase of computational costs. In this paper, we propose a different approach which does not need in adding points to Ω_s and which is cheaper than the use of the standard CEM on the extended uniform grid $\tilde{\Omega}_s$.

Consider the covariance matrix R of the random vector $Z(s_j^{(i)})$, $s_j^{(i)} \in \Omega_s$, written in the block matrix form:

$$R = \begin{bmatrix} R_{0,0} & R_{0,1} & R_{0,2} & \cdots & R_{0,N-1} \\ R_{1,0} & R_{1,1} & R_{1,2} & \cdots & R_{1,N-1} \\ R_{2,0} & R_{2,1} & R_{2,2} & \cdots & R_{2,N-1} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ R_{N-1,0} & R_{N-1,1} & R_{N-1,2} & \cdots & R_{N-1,N-1} \end{bmatrix}_{N\ell \times N\ell}, \quad (2.1)$$

where each block matrix $R_{i,k}$ is defined as

$$R_{i,k} = \left[r(|s_j^{(i)} - s_l^{(k)}|) \right]_{1 \leq j, l \leq \ell}. \quad (2.2)$$

Now note that, by construction,

$$R_{i,j} = \begin{cases} R_{k,l} & \text{if } j - i = l - k, \\ R_{k,l}^T & \text{if } j - i = k - l. \end{cases} \quad (2.3)$$

Property (2.3) implies that the covariance matrix R from (2.1) can be uniquely determined by its first block row and hence it is symmetric and block Toeplitz, having identical blocks along diagonals. Then R can be rewritten as

$$R = \begin{bmatrix} R_{0,0} & R_{0,1} & R_{0,2} & \cdots & R_{0,N-1} \\ R_{0,1}^T & R_{0,0} & R_{0,1} & \cdots & R_{0,N-2} \\ R_{0,2}^T & R_{0,1}^T & R_{0,0} & \cdots & R_{0,N-3} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ R_{0,N-1}^T & R_{0,N-2}^T & R_{0,N-3}^T & \cdots & R_{0,0} \end{bmatrix}. \quad (2.4)$$

We now illustrate how CEM [22, 6, 23] can be extended so that its new version, BCEM, is applicable to the symmetric block Toeplitz matrix R from (2.4). To this end, we embed R in the $m\ell \times m\ell$ symmetric block Toeplitz matrix C for some even integer $m \geq 2N$:

$$C = \begin{bmatrix} C_0 & C_1 & \cdots & C_{m-1} \\ C_{m-1} & C_0 & \cdots & C_{m-2} \\ \vdots & \vdots & \ddots & \vdots \\ C_1 & C_2 & \cdots & C_0 \end{bmatrix}, \quad (2.5)$$

where

$$C_k = \left[r(g(|s_i^{(0)} - s_j^{(k)}|)) \right]_{1 \leq i, j \leq \ell} \quad (2.6)$$

and

$$g(x) = \begin{cases} x & \text{if } x < mh/2, \\ mh - x & \text{if } x \geq mh/2. \end{cases} \quad (2.7)$$

Note that $C_i = R_{0,i}$ for $0 \leq i \leq N-1$ and that C is the covariance matrix for $Z(x)$ defined in the circular manner on the grid $\Omega_s^E = \bigcup_{i=0}^{m-1} S_i \subset \Omega^E = [x_0, x_m]$, where $x_m = x_0 + mh$ and S_i are defined in the same way as before.

It is not difficult to see that the matrix C has the following properties

$$C_0 = C_0^T, \quad (2.8)$$

$$C_k = C_{m-k}^T, \text{ for } 1 \leq k \leq \frac{m}{2}. \quad (2.9)$$

The properties (2.8) and (2.9) imply that C is a symmetric block circulant matrix.

Let F_B be the tensor product of a one-dimensional discrete Fourier matrix F_m^1 of order m and an identity matrix I_ℓ of size $\ell \times \ell$:

$$F_B = F_m^1 \otimes I_\ell = \begin{bmatrix} I_\ell & I_\ell & I_\ell & \cdots & I_\ell \\ \omega_0 I_\ell & \omega_1 I_\ell & \omega_2 I_\ell & \cdots & \omega_{m-1} I_\ell \\ \omega_0^2 I_\ell & \omega_1^2 I_\ell & \omega_2^2 I_\ell & \cdots & \omega_{m-1}^2 I_\ell \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \omega_0^{m-1} I_\ell & \omega_1^{m-1} I_\ell & \omega_2^{m-1} I_\ell & \cdots & \omega_{m-1}^{m-1} I_\ell \end{bmatrix}.$$

The matrix C is unitarily block diagonalizable by F_B [3, 23], i.e., there exists $\ell \times \ell$ matrices Λ_k , $k = 0, 1, \dots, m-1$, such that

$$C = \frac{1}{m} F_B \Lambda F_B^H, \quad \text{where } \Lambda = \begin{bmatrix} \Lambda_0 & 0 & \cdots & 0 \\ 0 & \Lambda_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \Lambda_{m-1} \end{bmatrix}. \quad (2.10)$$

Here H denotes the conjugate transpose. Similarly to the eigenvalue decomposition of a symmetric circulant matrix whose eigenvalues can be calculated by performing a discrete Fourier transform of its first row (or column), the block matrices on the diagonal of Λ can be computed as

$$[C_0 \ C_1 \ \cdots \ C_{m-1}] F_B = [\Lambda_0 \ \Lambda_1 \ \cdots \ \Lambda_{m-1}] \quad (2.11)$$

or in the component-wise form:

$$\begin{bmatrix} C_0^{i,j} & C_1^{i,j} & \cdots & C_{m-1}^{i,j} \end{bmatrix} F_m^1 = \begin{bmatrix} \Lambda_0^{i,j} & \Lambda_1^{i,j} & \cdots & \Lambda_{m-1}^{i,j} \end{bmatrix}, \text{ where } 1 \leq i, j \leq \ell. \quad (2.12)$$

Since the block circulant matrix C is real and symmetric, Λ_k are Hermitian. Furthermore, all the diagonal elements of Λ_k are equal. Therefore, only $\ell(\ell+1)/2 - (\ell-1)$ applications of F_m^1 are required for computing Λ .

Remark 2.1. Consider a uniform grid, i.e., a block-regular grid with the size of regular grid being a multiple of the number of blocks or, in other words, the points in each S_i being uniformly located. Then BCEM is applicable on the uniform grid (recall that CEM works on regular grids only). Since the covariance depends on the distance between points only, the block circulant matrix C on the uniform grid satisfies the relationship

$$C_k^{a,b} = C_k^{c,d} \text{ if } |a-b| = |c-d|.$$

Consequently, Λ_k in (2.11) are Toeplitz and the number of 1D FFT F_m^1 to compute distinctive values of Λ is equal to ℓ . Thus, in BCEM the block circulant matrix can be diagonalized by using ℓ FFTs of order m followed by using a Cholesky decomposition of the block diagonal matrix Λ , whose block entries are of size $\ell \times \ell$. For small ℓ , the overall computational cost is dominated by $\mathcal{O}(\ell m \log_2 m)$. On the other hand, the complexity of CEM is dominated by FFTs of order $m\ell$, which gives the overall cost $\mathcal{O}(m\ell \log_2 m\ell)$. Hence, BCEM can outperform CEM on the uniform grid, where both CEM and BCEM use the same covariance matrix (see also Remark 4.2).

The symmetricity of C also guaranties the spectral decomposition

$$\Lambda_k = U_k D_k U_k^H, \quad (2.13)$$

where U_k is unitary and D_k is a real-valued diagonal matrix. The following proposition implies that Λ from (2.10) can be decomposed with $m/2 + 1$ applications of the spectral decompositions (2.13).

Proposition 2.2. *The block diagonal matrix Λ from (2.10) has the property*

$$\Lambda_k = \overline{\Lambda_{m-k}}, \text{ for } 1 \leq k \leq \frac{m}{2}, \quad (2.14)$$

where the bar denotes the matrix with conjugate complex entries.

Proof. Let $\omega_n = \exp(\frac{2\pi n}{m}i)$ be a root of unity. Then (see (2.9) and (2.12)):

$$\begin{aligned} \overline{\Lambda_{m-k}} &= C_0 + \overline{\omega_{m-k}} C_1 + \cdots + \overline{\omega_{m-k}}^{m-1} C_{m-1} \\ &= C_0 + \omega_k C_1 + \cdots + \omega_k^{m-1} C_{m-1} \\ &= \Lambda_k. \end{aligned}$$

□

It follows from (2.10) and (2.13) that C has the eigenvalue decomposition

$$C = \frac{1}{m} (F_B U) D (F_B U)^H, \quad (2.15)$$

where the unitary block-diagonal matrix U and the diagonal matrix D are of the form

$$U = \begin{bmatrix} U_0 & 0 & \cdots & 0 \\ 0 & U_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & U_{m-1} \end{bmatrix} \text{ and } D = \begin{bmatrix} D_0 & 0 & \cdots & 0 \\ 0 & D_1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & D_{m-1} \end{bmatrix}.$$

We note that C is non-negative definite if and only if $D_k^{i,i} \geq 0$ for each $0 \leq k \leq m-1$ and $1 \leq i \leq \ell$.

Assume for the moment that all the eigenvalues of C are non-negative. Let two independent random vectors ξ_1 and ξ_2 , each of size m , be normally distributed $\mathcal{N}(O, I_m)$, i.e., $E[\xi_i \xi_j^T] = \delta_{ij} I_m$, where δ_{ij} denotes

the Kronecker delta. Set $\eta = U(D/m)^{\frac{1}{2}}(\xi_1 + i\xi_2)$. Then the real and imaginary parts of the vector $\zeta := F_B\eta$ give two independent random vectors ζ_1 and ζ_2 that are both distributed as $N(0, C)$. Since R is embedded in C , the corresponding parts of ζ_1 and ζ_2 are distributed as $\mathcal{N}(0, R)$. Note that the matrix-vector multiplication $F_B\eta$ can be calculated component-wise by ℓ applications of F_m^1 .

The algorithm described above depends on non-negative definiteness of the symmetric block circulant matrix C . The sufficient conditions for symmetric circulant matrices to have all non-negative eigenvalues were developed for 1D case in [6] and [22]. Here we extend these conditions to the symmetric block circulant matrix C from (2.5). To this end, introduce a uniform grid $\tilde{\Omega}_s$ such that $\Omega_s \subset \tilde{\Omega}_s$ and consider the covariance matrix \tilde{R} defined on $\tilde{\Omega}_s$. As Ω_s is a subset of $\tilde{\Omega}_s$, R is a sub-matrix of the matrix \tilde{R} . Let a uniform grid $\tilde{\Omega}_s^E$ contain all points of $\Omega_s^E = \bigcup_{i=0}^{m-1} S_i$. Then the symmetric block circulant matrix C is a sub-matrix of a symmetric circulant matrix \tilde{C} :

$$\tilde{C}^{i,j} = [r(g(|x_i - x_j|))], \quad (2.16)$$

where the function $g(x)$ is as in (2.7) and $x_i, x_j \in \tilde{\Omega}_s^E$. Therefore, there exists an injection matrix P^T such that

$$C = P^T \tilde{C} P. \quad (2.17)$$

An injection matrix can be built by eliminating rows of the identity matrix, which correspond to points not in Ω_s^E . For instance,

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

is an injection matrix from $\{x_1, x_2, x_3, x_4, x_5\}$ to $\{x_1, x_3, x_5\}$. The relationship (2.17) leads to the following proposition.

Proposition 2.3. *If \tilde{C} is non-negative definite, then so is C .*

When the circulant matrix \tilde{C} fails to be non-negative definite, Wood and Chen [22] suggested to increase the size of \tilde{C} until it becomes non-negative definite (the so-called padding technique). From the relationship (2.17) between C and \tilde{C} , the same strategy can be used for the matrix C . That is, increase m until C becomes non-negative definite. Therefore, the number of blocks m , which is required for C to be non-negative definite, depends on the grid size of the uniform grid $\tilde{\Omega}_s^E$, not on the number of points in Ω_s . Thus, the number of paddings needed for BCEM is the same as for CEM (see also Remark 3.2).

3 Multidimensional BCEM

In the previous section we illustrated the idea of BCEM in the simpler setting of 1D space. In this section we present multi-dimensional BCEM which computational complexity is discussed in the next section.

We start with introducing the notation which largely follows [23]. Let \mathbb{Z}^d be the set of d -vectors with non-negative integer components and $\mathbf{0}$ and $\mathbf{1}$ be the d -dimensional vectors whose all components equal to 0 and 1, respectively. For $\mathbf{i} = (\mathbf{i}[1], \dots, \mathbf{i}[d])^T$, $\mathbf{j} = (\mathbf{j}[1], \dots, \mathbf{j}[d])^T \in \mathbb{Z}^d$, we define addition in \mathbb{Z}^d :

$$\mathbf{i} + \mathbf{j} = (\mathbf{i}[1] + \mathbf{j}[1], \dots, \mathbf{i}[d] + \mathbf{j}[d])^T$$

and also the product of elements of \mathbf{i} :

$$\bar{\mathbf{i}} = \prod_{k=1}^d \mathbf{i}[k].$$

For any $\mathbf{j} \in \mathbb{Z}^d$ all components of which are strictly positive, we define the set $\mathcal{I}(\mathbf{j})$:

$$\mathcal{I}(\mathbf{j}) = \{\mathbf{i} \in \mathbb{Z}^d : 0 \leq \mathbf{i}[k] \leq \mathbf{j}[k] - 1 \text{ for all } 1 \leq k \leq d\}. \quad (3.1)$$




Figure 3.1: The 2D uniform grid $\Omega_r^E = \{x_{\mathbf{j}_0}=x_{\mathbf{0}}, x_{\mathbf{j}_1}, \dots, x_{\mathbf{j}_{\mathbf{m}+1}-1}=x_{\mathbf{m}}\}$ on the rectangle $\Omega^E = [x_{\mathbf{0}}[1], x_{\mathbf{m}}[1]] \times [x_{\mathbf{0}}[2], x_{\mathbf{m}}[2]]$ for $\mathbf{m} = (8, 8)^T$. Nodes $s_{\mathbf{j}}^{(\mathbf{j}_k)}$ of the block-regular grid Ω_s^E are represented by black circles. The shaded rectangle corresponds to the computation (i.e., the domain of interest for the problem at hands) domain $\Omega = [x_{\mathbf{0}}[1], x_{\mathbf{N}}[1]] \times [x_{\mathbf{0}}[2], x_{\mathbf{N}}[2]]$ for $\mathbf{N} = (4, 4)^T$.

Note that the cardinality of $\mathcal{I}(\mathbf{j})$ is equal to $\bar{\mathbf{j}}$.

Introduce a d -dimensional rectangular parallelepiped

$$\Omega = \prod_{i=1}^d [x_{\mathbf{0}}[i], x_{\mathbf{N}}[i]] \subset \mathbb{R}^d, \quad (3.2)$$

where $\mathbf{N} = (\mathbf{N}[1], \dots, \mathbf{N}[d])^T \in \mathbb{Z}^d$ and the vector \mathbf{h} with the components $\mathbf{h}[i] = (x_{\mathbf{N}}[i] - x_{\mathbf{0}}[i]) / \mathbf{N}[i]$. Further,

points $x_{\mathbf{i}_k} = (x_{\mathbf{i}_k}[1], \dots, x_{\mathbf{i}_k}[d])^T$ with $x_{\mathbf{i}_k}[j] = x_{\mathbf{0}}[j] + \mathbf{i}_k[j]\mathbf{h}[j]$ form a regular grid $\Omega_r = \{x_{\mathbf{i}_0}, \dots, x_{\mathbf{i}_{\mathbf{N}+1}-1}\}$ on the rectangular parallelepiped Ω (see Fig. 3.1). The domain Ω in (3.2) can be divided into d -dimensional rectangular parallelepipeds as

$$\Omega = \bigcup_{\mathbf{j}_k \in \mathcal{I}(\mathbf{N})} \Omega_{\mathbf{j}_k}, \quad (3.3)$$

where

$$\Omega_{\mathbf{j}_k} = \prod_{i=1}^d [x_{\mathbf{j}_k}[i], x_{\mathbf{j}_k}[i] + \mathbf{h}[i]], \mathbf{j}_k \in \mathcal{I}(\mathbf{N}). \quad (3.4)$$

For the purpose of algorithm development, we use a lexicographic ordering of \mathbf{j}_k with respect to k , i.e., row after row and layer after layer (see Fig. 3.1).

Consider a stationary Gaussian random field $Z(x)$, $x \in \mathbb{R}^d$, with zero mean and covariance function $r(x)$. We assume that the problem at our hands is such that we need to sample $Z(x)$ at the nodes $s_i^{(\mathbf{j}_k)}$ defined as follows (see the examples in the Introduction and also Fig. 3.1):

$$s_j^{(\mathbf{j}_k)} = x_{\mathbf{j}_k} + \delta_j, \quad (3.5)$$

where $\delta_j = (\delta_j[1], \dots, \delta_j[d])^T$ with $0 \leq \delta_j[i] < \mathbf{h}[i]$ for $1 \leq i \leq d$. Here ℓ is the number of sampling points in each subdomain $\Omega_{\mathbf{j}_k}$. That is, in each $\Omega_{\mathbf{j}_k}$ the points from the set $S_{\mathbf{j}_k} = \{s_1^{(\mathbf{j}_k)}, \dots, s_\ell^{(\mathbf{j}_k)}\}$ are distributed according to the same pattern for all $\mathbf{j}_k \in \mathcal{I}(\mathbf{N})$. Denote the grid:

$$\Omega_s = \bigcup_{\mathbf{j}_k \in \mathcal{I}(\mathbf{N})} \Omega_{\mathbf{j}_k}.$$

Note that δ_j are independent of the index vector \mathbf{j}_k . The covariance matrix R of $Z(s_i^{(\mathbf{j}_k)})$, $s_i^{(\mathbf{j}_k)} \in \Omega_s$, is block-Toeplitz. In the one-dimensional-case it consists of non-Toeplitz blocks of order ℓ (see Section 2). In the d -dimensional case with $d > 1$, it consists of blocks which have all the properties of a correlation matrix in the $d - 1$ dimensional space. We emphasize that the matrix R is not Toeplitz and hence CEM is not directly applicable here.

Analogously to CEM, in order to build a block-circulant matrix, we consider an extended domain $\Omega^E = \prod_{i=1}^d [x_{\mathbf{0}}[i], x_{\mathbf{m}}[i]]$, where $\mathbf{m} = (\mathbf{m}[1], \dots, \mathbf{m}[d])^T$ with $\mathbf{m}[i] \geq 2\mathbf{N}[i]$ and $x_{\mathbf{m}}[i] = x_{\mathbf{0}}[i] + \mathbf{m}[i]\mathbf{h}[i]$ for $1 \leq i \leq d$. Figure 3.1 shows an example of the computation domain Ω with $\mathbf{N} = (4, 4)^T$ and the extended domain Ω^E with $\mathbf{m} = (8, 8)^T$. Vectors $\mathbf{j}_k \in \mathcal{I}(\mathbf{m} + \mathbf{1})$ form the extended regular grid $\Omega_r^E = \{x_{\mathbf{j}_k} = (x_{\mathbf{j}_k}[1], \dots, x_{\mathbf{j}_k}[d])^T \mid \mathbf{j}_k \in \mathcal{I}(\mathbf{m} + \mathbf{1})\} \subset \Omega^E$. There are $\mathbf{m} + \mathbf{1}$ regular grid points in the set Ω_r^E . The parallelepiped Ω^E can be divided into d -dimensional small parallelepipeds as (see also Fig. 3.1):

$$\Omega^E = \bigcup_{\mathbf{j}_k \in \mathcal{I}(\mathbf{m})} \Omega_{\mathbf{j}_k}, \quad (3.6)$$

where $\Omega_{\mathbf{j}_k}$, $\mathbf{j}_k \in \mathcal{I}(\mathbf{m})$, are as in (3.4).

We now describe BCEM in the d -dimensional case, which is applicable to our block-Toeplitz covariance matrix R . In contrast to the 1D setting, where the covariance function is always even because of its symmetry, further classifications of covariance functions are needed in higher dimensional cases. We say that r is *component-wise even* in i th coordinate if

$$r(x[1], \dots, -x[i], \dots, x[d]) = r(x[1], \dots, x[i], \dots, x[d]) \quad (3.7)$$

for all $x \in \mathbf{R}^d$; otherwise, we say that r is *component-wise uneven* in some coordinates.

For simplicity of the exposition, let us assume for now that $r(x)$ is component-wise even in all coordinates (we will discuss a modification of BCEM in the uneven case in Remark 3.1).

We first build the block circulant embedding of the block-Toeplitz matrix R . Consider the first row of the block circulant matrix C , which is an $\ell \times \ell \bar{\mathbf{m}}$ matrix C_f of the form

$$C_f = [C_{\mathbf{j}_0} \quad C_{\mathbf{j}_1} \quad \cdots \quad C_{\mathbf{j}_{\bar{\mathbf{m}}-1}}], \quad (3.8)$$

where (i, j) -th element of $C_{\mathbf{j}_k}$ is equal to

$$C_{\mathbf{j}_k}^{i,j} = r(g_{\mathbf{m}}(s_i^{(\mathbf{j}_0)} - s_j^{(\mathbf{j}_k)})) \quad (3.9)$$

and the vector function $g_{\mathbf{m}} = (g_{\mathbf{m}}^1, \dots, g_{\mathbf{m}}^d)^T$ is defined by

$$g_{\mathbf{m}}^i(x_{\mathbf{j}}) = \begin{cases} x_{\mathbf{j}}[i], & \text{if } |x_{\mathbf{j}}[i]| < \mathbf{m}[i]\mathbf{h}[i]/2, \\ \mathbf{m}[i]\mathbf{h}[i] - |x_{\mathbf{j}}[i]|, & \text{otherwise.} \end{cases} \quad (3.10)$$

The block circulant matrix C is generated by its first row C_f in the usual way. Also note that

$$R = [C_{\mathbf{j}_k}]_{\mathbf{j}_k \in \mathcal{I}(\mathbf{N})}. \quad (3.11)$$

The block circulant matrix C is block diagonalizable by a block discrete Fourier transform (BDFT) matrix, $F_B = F_{\mathbf{m}}^d \otimes I_\ell$, where $F_{\mathbf{m}}^d$ is a d -dimensional DFT matrix. That is, we have $C = (1/\bar{\mathbf{m}})F_B \Lambda F_B^H$, where $\Lambda = \text{diag}(\Lambda_0, \dots, \Lambda_{\bar{\mathbf{m}}})$. The blocks on the diagonal of Λ can be found by simply taking BDFT of first block row [3, 23]. Furthermore, using the fact that F_B is the tensor product involving the identity matrix, we derive the following component-wise computation:

$$[\Lambda_0^{i,j} \dots \Lambda_{\bar{\mathbf{m}}-1}^{i,j}] = \text{FFT}_d([C_0^{i,j} \dots C_{\bar{\mathbf{m}}-1}^{i,j}]), \quad (3.12)$$

where $1 \leq i, j \leq \ell$ and FFT_d is the d -dimensional FFT. Note that instead of FFT_1 we will write FFT.

Due to the fact that Λ is Hermitian and all diagonal entries of Λ_k are the same, the required number of FFT_d of size \mathbf{m} (which is equivalent to FFT of order $\bar{\mathbf{m}}$) in (3.12) is $\ell(\ell+1)/2 - (\ell-1)$.

If Λ is positive-definite, the Cholesky decomposition $\Lambda = LL^H$ exists, where L is a block diagonal matrix with each block being a lower triangular matrix. Then, we obtain the decomposition $C = (1/\bar{\mathbf{m}})F_B L (F_B L)^H$.

As in the one-dimensional case (see Section 2), let $\xi = \xi_1 + i\xi_2$ be a complex-valued random vector of order $\bar{\mathbf{m}}$ with ξ_1 and ξ_2 being real, normal random vectors such that $E[\xi_i] = 0$ and $E[\xi_i \xi_j^T] = \delta_{ij}I$. Set $\tilde{L} = (1/\bar{\mathbf{m}})^{1/2}L$ and $\eta = \tilde{L}\xi$. Multiplying the square root of C by ξ , we obtain the complex-valued vector

$$F_B \eta = \zeta = \zeta_1 + i\zeta_2, \quad (3.13)$$

with the properties: $E[\zeta_1 \zeta_1^T] = E[\zeta_2 \zeta_2^T] = C$ and ζ_1 and ζ_2 are independent. Using tensor-product properties of F_B , ζ can be computed in the component-wise manner:

$$[\zeta[i] \quad \zeta[i+\ell] \dots \zeta[i+(\bar{\mathbf{m}}-1)\ell]] = \text{FFT}_d([\eta[i] \quad \eta[i+\ell] \dots \eta[i+(\bar{\mathbf{m}}-1)\ell]]) \quad (3.14)$$

for $1 \leq i \leq \ell$.

To summarize, the new BCEM can be presented in the algorithmic form as follows.

Algorithm 3.1 Block circulant embedding method (BCEM)

Given $N \in \mathbb{Z}^d$, $x_0 \in \Omega$, and strictly positive valued vector $\mathbf{h} \in \mathbf{R}^d$,

- Step 1.* Choose a vector $\mathbf{m} \in \mathbb{Z}^d$ such that $\mathbf{m}[i] \geq 2N[i]$ for all $1 \leq i \leq d$.
- Step 2.* Compute the first block row of the circulant matrix C as described in (3.8)-(3.10).
- Step 3.* Compute the block diagonal matrix $\Lambda = \text{diag}(\Lambda_0, \dots, \Lambda_{\overline{\mathbf{m}}-1})$ using (3.12).
- Step 4.* Compute the square-root of Λ applying Cholesky decompositions to diagonal blocks of Λ :

$$\Lambda = LL^H, \quad (3.15)$$

where L is a block diagonal matrix with lower triangular block of order ℓ .

Step 5. If the Cholesky decomposition fails in *Step 4*, increase $\mathbf{m}[i]$ by one or more and go to *Step 2*.

Step 6. Compute $\tilde{L} = (1/\overline{\mathbf{m}})^{1/2}L$.

Step 7. Generate a random complex vector of dimension $\overline{\mathbf{m}}\ell$, $\xi = \xi_1 + i\xi_2$, with two independent vectors ξ_1 and ξ_2 being $\mathcal{N}(0, I_{\overline{\mathbf{m}}\ell})$. Compute $\eta = \tilde{L}\xi$.

Step 8. Compute $z = (\zeta[1], \dots, \zeta[\overline{\mathbf{m}}\ell])^T$ by applying times as in (3.14).

Note that if $\ell = 1$, then Ω_r is regular, C is circular, Λ becomes diagonal instead of block diagonal and Algorithm 3.1 degenerates to the standard CEM.

Remark 3.1. Algorithm 3.1 is applicable when the covariance function is component-wise even (see (3.7)). Although the covariance function is even by definition, i.e., $r(-x) = r(x)$, it could be component-wise uneven, e.g.,

$$r(x) = \exp(-x^T A x) \quad \text{with} \quad A = \begin{bmatrix} 2 & -1 \\ -1 & 2 \end{bmatrix}, \quad x \in \mathbb{R}^2,$$

is uneven. In this case, the matrix C defined by the vector function g_m in (3.10) usually does not have block circulant structure because of conflicting definitions at the points x with $x[i] = \mathbf{m}[i]\mathbf{h}[i]/2$, if r is uneven in i th coordinate. Two adjustments to make CEM work in uneven cases were suggested in [22], which can be applied to BCEM by modifying Algorithm 3.1 as follows.

If r is uneven in the i th coordinate, either

- (a) choose $\mathbf{m}[i]$ to be an odd integer, e.g., a power of three;

or

- (b) still choose $\mathbf{m}[i]$ to be an even integer and define C_{j_k} using (3.9) and (3.10), except put $r(x) = 0$ for all x such that $|x[i]| = \mathbf{m}[i]\mathbf{h}[i]/2$ for some i .

In either case, the resulting matrix C has a block circulant structure and thus Algorithm 3.1 can be seamlessly extended to the uneven case with aforementioned modifications.

Remark 3.2. As mentioned earlier, the matrix C is often negative definite in practical applications of CEM and BCEM. Following [22], we increase the matrix C in Algorithm 3.1 (see its Step 5) until it becomes non-negative definite (the padding technique). The padding technique is universal and usually efficient when the correlation length of a random field is in a range from small to medium relative to the size of a computational domain and the field is not too smooth. Otherwise, the use of the padding technique could be very expensive. There are two recently developed alternatives to padding (a cut-off of the circulant matrix [19, 12] and smoothing window circulant embedding [14]), which can deal with the problem of negative definiteness of circulant matrices effectively. The techniques from [19, 12, 14] are applicable to BCEM as they are for CEM.

The equispaced FFT is highly parallelizable in high dimensions, and its highly scalable implementations are proposed in [8, 7]. This could be beneficial in the standard CEM because its computation is dominated by the FFT. Still equipped with the parallelism of the FFT, BCEM can be further parallelized in a natural

way because the applications of FFT_d in *Step 3* and *Step 8* of Algorithm 3.1 can be performed separately and simultaneously. Moreover, block-diagonal matrix operations in *Step 4* and *Step 7* can be performed separately and simultaneously. Therefore, the overall BCEM algorithm contains two-level parallelism, giving us significant advantage over the standard CEM.

As we will see in the next section, BCEM can be faster than CEM both in taking square-roots of the corresponding circulant matrices (performed, of course, only once per the whole Monte Carlo simulation) and in sampling the random field required in each Monte Carlo run. The latter is usually more important in Monte Carlo-type simulations.

4 Computational Complexity of BCEM

In this section we analyze the computational complexity of BCEM. To this end, we use the same convention as in Golub and Van Loan [13] for counting the number of floating point operations: $5m \log_2 m$ flops for FFT of size m and $n^3/3$ flops for the Cholesky decomposition of a matrix of order n .

Step 3 of Algorithm 3.1 is the initial factorization of the block circulant matrix C by taking BDFT of its first block row which can be computed using the ordinary DFT in (3.12) at the cost

$$\text{cost}_1 = \left(\frac{\ell(\ell+1)}{2} - (\ell-1) \right) (5\overline{\mathbf{m}} \log_2 \overline{\mathbf{m}}) \text{ flops.} \quad (4.1)$$

Here we took into account that each Λ_k is Hermitian and its diagonal elements have the same value.

In *Step 4*, the square-root operation on the block diagonal matrix Λ with $\overline{\mathbf{m}}$ blocks of order ℓ can be performed on each block separately using the Cholesky decomposition method. In Proposition 2.2, we proved in the one-dimensional case that Λ has pairs of complex-conjugate blocks, Λ_k and Λ_{m-k} , which allows us to compute the square-root of Λ_k and use its complex-conjugate as a square-root of its complex-conjugate pair Λ_{m-k} . This is based on the periodicity and conjugate symmetry of FFT. Hence, Proposition 2.2 can be extended to the higher dimensional cases. Then the matrix Λ can be decomposed at the cost

$$\text{cost}_2 = \prod_{i=1}^d \left(\frac{\mathbf{m}[i]}{2} + 1 \right) \frac{\ell^3}{3} \text{ flops.} \quad (4.2)$$

Remark 4.1. Note that if the nodes of $S_{\mathbf{j}_k}$ are regularly (uniformly) distributed in $\Omega_{\mathbf{j}_k}$ for all $\mathbf{j}_k \in I(\mathbf{m})$ which is often the case in applications (see, e.g., Example 1.1), then all blocks on the diagonal of Λ are block Toeplitz (Toeplitz). Toeplitz matrix and block Toeplitz matrix can be decomposed using Schur's algorithm [20] and block Schur's algorithm [9], respectively, which have $O(\ell^2)$ complexity as opposed to $O(\ell^3)$ for the standard Cholesky decomposition. Making use of Schur's algorithms can reduce the cost of Algorithm 3.1.

In *Step 8*, computing a realization of ζ requires block diagonal matrix-vector multiplication $\tilde{L}\xi$ and ℓ applications of FFT of order $\overline{\mathbf{m}}$ in (3.14) at the cost

$$\text{cost}_3 = \ell^2 \overline{\mathbf{m}} \text{ flops} \quad (4.3)$$

and

$$\text{cost}_4 = \ell(5\overline{\mathbf{m}} \log_2 \overline{\mathbf{m}}) \text{ flops,} \quad (4.4)$$

respectively.

To conclude, the cost of BCEM is $\mathcal{O}(\ell^3 \overline{\mathbf{m}} + \ell^2 \overline{\mathbf{m}} \log_2 \overline{\mathbf{m}})$ flops. In practical applications of BCEM (see, e.g. examples in the Introduction) the size of blocks ℓ is relatively small while the number of blocks $\overline{\mathbf{m}}$ is large. Recall that BCEM is designed for block-regular grids Ω_s . Its main computational advantage in comparison with CEM (which is designed for regular grids) comes from the fact that the use of CEM in the case of simulations on a block-regular grid Ω_s requires regularization of Ω_s , i.e., adding a significant number of extra nodes which BCEM does not need. Hence BCEM works on a grid with a smaller number of nodes than CEM and needs to generate random vectors ζ of smaller size than CEM (and hence makes less number of calls to a random number generator to sample ξ).

Remark 4.2. It can be shown that the use of BCEM on a *regular* grid split in blocks of a size ℓ can be more effective in sampling the random field than CEM but it is computationally more expensive in the matrix decomposition than CEM. The latter can be overcome by exploiting the fact that BCEM is parallelizable in comparison with CEM. Thus BCEM can be more effective than CEM even in the case of regular grids for which CEM is designed.

We now compare computational complexity of BCEM and CEM using the first two examples from the Introduction and the following exponential covariance function (cf. (3.9)):

$$r(\mathbf{x}) = \sigma^2 \exp\left(-\frac{\|\mathbf{x}\|_1}{0.3}\right), \quad (4.5)$$

where $\|\cdot\|_1$ means L_1 -norm. We note that the circulant matrix C (cf. (3.8), (3.9)) of the size $\mathbf{m} = 2\mathbf{N}$ formed by (4.5) is always positive definite (see, e.g. [6]). This means, in particular, that Step 5 (i.e., padding) of Algorithm 3.1 is not needed in this case. For simplicity, we consider the domain Ω to be the unit square in the examples.

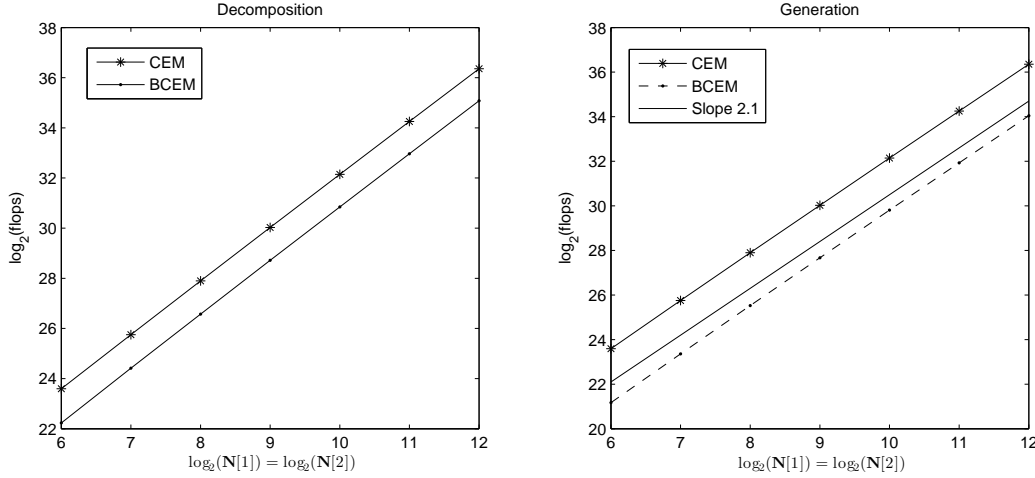



Figure 4.1: The floating point operations required for the matrix decomposition and random field generation stages of CEM and BCEM in Example 4.3.

Example 4.3. *Triangular finite element with a quadrature point located at the barycentre of the triangle.*

In Example 1.1 (see Figure 1.1), each rectangular block contains 9 uniform grid nodes. Hence the order of the circulant matrix used by CEM is $9\overline{\mathbf{m}}$, where $\mathbf{m} = 2\mathbf{N}$ and $\overline{\mathbf{m}}$ is the number of rectangular blocks in the extended domain Ω^E . Then the matrix decomposition cost for CEM is $45\overline{\mathbf{m}} \log_2 9\overline{\mathbf{m}}$ flops, and generation of each realization of the random field requires another $45\overline{\mathbf{m}} \log_2 9\overline{\mathbf{m}}$ flops.

Here BCEM uses only two points in each rectangular block, so the order of the block-circulant matrix is $2\overline{\mathbf{m}}$. Substituting $\ell = 2$ into (4.1) and (4.2), the total matrix decomposition cost for BCEM is $10\overline{\mathbf{m}} \log_2 \overline{\mathbf{m}} + (\mathbf{m}[1]/2 + 1)(\mathbf{m}[2]/2 + 1)(8/3)$ flops. Each realization of the random field is generated at the cost of $4\overline{\mathbf{m}} + 10\overline{\mathbf{m}} \log_2 \overline{\mathbf{m}}$ flops (see (4.3) and (4.4)).

Figure 4.1 shows the floating point operations required by the two algorithms. One can see that BCEM is more effective in both procedures and that the complexity of BCEM grows at roughly the same rate as for CEM. Compared to CEM, BCEM reduces the matrix decomposition cost and the generation cost approximately in 2.5 time and 4 time, respectively. The improvement in computational efficiency is due to the fact that BCEM works with just 2/9 of nodes that CEM uses to build the circulant matrix. This also means that BCEM requires 4.5 time less memory than CEM.



N	CPU Time (s)		
	BCCEM	CEM	speed-up
32	0.60	2.71	4.5
64	2.32	11.39	4.9
128	11.34	48.55	4.3
256	49.37	212.21	4.3

Figure 4.2 & Table 4.1: Average time required to simulate a single realization of the random field in Example 4.3 computed using 1000 independent samples.

To compare performance of BCCEM and CEM further, we generated samples of the random field by these two methods on an Intel Xeon E5-2450, 96GB RAM computer using MATLAB R2014a. Figure 4.2 shows the average computational time of generation of a single realization of the random field by both methods and how it increases with increasing N . Table 4.1 gives the CPU time and the speed up in generating a random vector using BCCEM against the ones using CEM. For both methods, the CPU time increases with increase of N at about the same rate as the theoretical rate shown in Figure 4.1 (right). We see that BCCEM is about 4.3 – 4.5 faster than CEM, which is close to the theoretical cost estimation in Figure 4.1.





Figure 4.3: The floating point operations required for the matrix decomposition and random field generation stages of CEM and BCCEM in Example 4.4.

Example 4.4. *Cell-centered finite volume discretization in multilevel Monte Carlo (MLMC) computation.* In Example 1.2, BCCEM uses 5 out of 16 uniform nodes required for CEM in each individual block to generate random variables located at the centers of both the fine and coarse cells. That is, CEM should generate random variables at the extra 11 nodes that are not used in the finite volume discretization and are not used by BCCEM. Then memory requirement for CEM and BCCEM are $16\overline{m}$ and $5\overline{m}$, respectively, which makes BCCEM more attractive when the number of blocks is large.

Whereas the matrix decomposition and sampling costs in CEM both require $80\bar{\mathbf{m}} \log_2 16\bar{\mathbf{m}}$ flops, the computational costs of the matrix decomposition and sampling in BCEM are $55\bar{\mathbf{m}} \log_2 \bar{\mathbf{m}} + (\mathbf{m}[1]/2 + 1)(\mathbf{m}[2]/2 + 1)(125/3)$ flops and $25\bar{\mathbf{m}} + 25\bar{\mathbf{m}} \log_2 \bar{\mathbf{m}}$ flops, respectively (see (4.1) - (4.4) with $\ell = 5$). Note that the total costs are dominated by $\bar{\mathbf{m}} \log_2 \bar{\mathbf{m}}$. Hence, for large $\bar{\mathbf{m}}$, the ratio of the matrix decomposition in CEM to one in BCEM is close to $80/55 \approx 1.45$. For the sample generation cost, the ratio is close to $80/25 \approx 3.2$. These theoretical computational costs are shown in Figure 4.3.



N	CPU Time (s)		
	CEM	BCEM	speed-up
32	0.94	3.08	3.3
64	4.62	13.71	3.0
128	23.37	68.28	2.9
256	101.82	301.24	3.0

Figure 4.4 & Table 4.2: Average time required to simulate a single realization of the random field in Example 4.4 computed using 1000 independent samples.

Figure 4.4 gives the CPU times for the random field generation. We see that the actual computational cost increases with increase of \mathbf{N} similarly to the theoretical one as in Figure 4.3. Table 4.2 demonstrates that BCEM is nearly 3 time faster than CEM as we expected from Table 1.1 and Figure 4.3. Also note that BCEM is highly parallelizable, so the computation cost can be further reduced using parallel algorithms.

We have compared BCEM and CEM on the 2D examples here. It is not difficult to see (cf. Table 1.1) that in 3D cases BCEM can outperform CEM even more dramatically.

Remark 4.5. The MATLAB codes for BCEM used for Examples 4.3 and 4.4 are available at <https://www.maths.nottingham.ac.uk/personal/pmzmp/bcem.html>.

Acknowledgments

This work was partially supported by the EPSRC grant EP/K031430/1.

References

- [1] A. BARTH, C. SCHWAB, AND N. ZOLLINGER, Multi-level Monte Carlo finite element method for elliptic PDEs with stochastic coefficients, *Numerische Mathematik*, 119 (2011), pp. 123–161.
- [2] K.A. CLIFFE, M.B. GILES, R. SCHEICHL, AND A.L. TECKENTRUP, Multilevel Monte Carlo methods and applications to elliptic PDEs with random coefficients, *Comput. Visual. Sci.*, 14 (2011), pp. 3–15.
- [3] P.J. DAVIS, Circulant Matrices, AMS Chelsea Publishing, 2nd ed., 1979.
- [4] J.P. DELHOMME, Spatial variability and uncertainty in groundwater flow parameters, a geostatistical approach, *Water Resour. Res.*, (1979), pp. 269–280.

- [5] C.R. DIETRICH AND G.N. NEWSAM, A fast and exact method for multidimensional Gaussian stochastic simulations: Extension to realizations conditioned on direct and indirect measurements, *Water Resour. Res.*, 32(6) (1996), pp. 1643–1652.
- [6] ———, Fast and exact simulation of stationary Gaussian processes through circulant embedding of the covariance matrix, *SIAM J. SCI. COMPUT.*, 18 (1997), pp. 1088–1107.
- [7] M. ELEFThERIOU, B. FITCH, A. RAYSHUBSKIY, T. WARD, AND R. GERMAIN, Performance measurements of the 3D FFT on the Blue Gene/L supercomputer, in *In Euro-Par 2005 Parallel Processing: 11th International Euro-Par Conference, Lisbon, Portugal, August 30-September 2, 2005*, J. Cunha and P. Medeiros, eds., vol. 3648 of *Lecture Notes in Computer Science*, Springer-Verlag, 2005, pp. 795–803.
- [8] M. FRIGO AND S. G. JOHNSON, FFTW: An adaptive software architecture for the FFT, in *In Proceedings of the IEEE International Conference on Acoustics*, vol. 3 of *Speech and Signal Processing (ICASSP)*, 1998, pp. 1381–1384.
- [9] K.A. GALLIVAN AND S. THIRUMALAI, High performance algorithm for Toeplitz and block Toeplitz matrices, *Linear Algebra and its Applications*, 241-243 (1996), pp. 343–388.
- [10] Y. GEL, A.E. RAFTERY, T. GNEITING, C. TEBALDI, D. NYCHKA, W. BRIGGS, M.S. ROULSTON, AND V.J. BERROCAL, Calibrated probabilistic mesoscale weather field forecasting: The geostatistical output perturbation method, *Journal of the American Statistical Association*, 99 (2004), pp. 575–590.
- [11] R.G. GHANEM AND P.D. SPANOS, Stochastic finite elements: A spectral approach, Springer, New York, 1991.
- [12] T. GNEITING, H. ŠEVČÍKOVÁ, D. B. PERCIVAL, M. SCHLATHER, AND JIAN Y., Fast and exact simulation of large Gaussian lattice systems in \mathbb{R}^2 : Exploring the limits, *J. Comput. Graph. Stat.*, 15 (2006), pp. 1–19.
- [13] G.H. GOLUB AND C.F. VAN LOAN, Matrix Computations, John Hopkins University Press, Baltimore, 3rd ed., 1996.
- [14] H. HELGASON, V. PIPIRAS, AND P. ABRY, Smoothing windows for the synthesis of gaussian stationary random fields using circulant matrix embedding, *J. Comput. Graph. Stat.*, 23 (2014), pp. 616–635.
- [15] O. LE MAÎTRE AND O. KNIO, Spectral Methods for Uncertainty Quantification With Applications to Computational Fluid Dynamics, Springer, 2010.
- [16] N. MARHEINEKE AND R. WEGNER, Fiber dynamics in turbulent flows: General modeling framework, *SIAM J. APPL. MATH.*, 66 (2006), pp. 1703–1726.
- [17] M. PARK AND K.A. CLIFFE, Conditional multilevel Monte Carlo simulation of groundwater flow in the Culebra Dolomite at the Waste Isolation Pilot Plant (WIPP) site, arXiv:1402.5257, (2014).
- [18] A.A. SKORDOS AND M.P.F. SUTCLIFFE, Stochastic simulation of woven composites forming, *Composites and Science and Technology*, 68 (2008), pp. 283–296.
- [19] M. STEIN, Fast and exact simulation of fractional brownian surfaces, *J. Comput. Graph. Stat.*, 11 (2002), p. 587599.
- [20] M. STEWART, Cholesky factorization of semidefinite Toeplitz matrices, *Linear Algebra and its Applications*, 254 (1997), pp. 497–525.
- [21] B. VERLEYE, D. NUYENS, A. WALBRAN, AND M. GAN, Uncertainty quantification in liquid composite moulding processes, in *Proceedings of the FPCM11 conference*, Aukland, July 2012, pp. 265–271.

- [22] A.T.A. WOOD AND G. CHAN, Simulation of stationary Gaussian processes in $[0, 1]^d$, Journal of Computational and Graphical Statistics, 3 (1994), pp. 409–432.
- [23] ———, Simulation of stationary Gaussian vector fields, Statistics and Computing, 9 (1999), pp. 265–268.
- [24] F. ZHANG, B. COSSON, S. COMAS-CARDONA, AND C. BINETRUY, Efficient stochastic simulation approach for RTM process with random fibrous permeability, Composites and Science and Technology, 71 (2010), pp. 1478–1485.